

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets

(11) Veröffentlichungsnummer:

(11) Publication number:

(11) Numéro de publication:

EP 1 527 645 A0

Internationale Anmeldung veröffentlicht durch die
Weltorganisation für geistiges Eigentum unter der Nummer:

WO 2020/040159 (art. 158 des EPÜ).

International application published by the World
Intellectual Property Organisation under number:

WO 2020/040159 (art. 158 of the EPC).

Demande internationale publiée par l'Organisation
Mondiale de la Propriété sous le numéro:

WO 2020/040159 (art. 158 de la CBE).

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
19 February 2004 (19.02.2004)

PCT

(10) International Publication Number
WO 2004/015904 A2

- (51) International Patent Classification⁷: **H04L**
- (21) International Application Number:
PCT/US2003/023942
- (22) International Filing Date: 31 July 2003 (31.07.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/214,997 8 August 2002 (08.08.2002) US
- (71) Applicant: HARRIS CORPORATION [US/US]; 1025
W. NASA Blvd., Melbourne, FL 32919 (US).
- (72) Inventor: CAIN, Joseph, Bibb; 200 Melbourne Avenue,
Indianapolis, IN 32903 (US).
- (74) Agents: YATSKO, Michael, S. et al.; Harris Corporation,
1025 W. Nasa Blvd., Melbourne, FL 32919 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU,
ZA, ZM, ZW.


(84) Designated States (regional): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,
ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO,
SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM,
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished
upon receipt of that report

[Continued on next page]

(54) Title: MULTIPLE PATH REACTIVE ROUTING IN A MOBILE AD HOC NETWORK



(57) Abstract: The method includes routing message data from a source node to a destination node in a mobile ad hoc network. The mobile ad-hoc network includes a plurality of intermediate mobile nodes between the source node and the destination node, and a plurality of wireless communication links connecting the nodes together. The method includes, at the source node, discovering routing to the destination node, ranking discovered routes according to at least one link metric, and distributing the message data to the destination node along a plurality of the discovered routes based upon the ranking. The link metric may include a measurement of link delay, link capacity, link available capacity, and/or link reliability.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

MULTIPLE PATH REACTIVE ROUTING IN A MOBILE AD HOC NETWORK**Background of the Invention**

A rapidly developing area of wireless networks is mobile ad hoc networks. Physically, a mobile ad hoc network includes a number of geographically-distributed, potentially mobile nodes wirelessly connected by one or more radio frequency channels. Compared with other type of networks, such as cellular networks or satellite networks, the most distinctive feature of mobile ad hoc networks is the lack of any fixed infrastructure. The network is formed of mobile nodes only, and a network is created on the fly as the nodes transmit to or receive from other nodes. The network does not in general depend on a particular node and dynamically adjusts as some nodes join or others leave the network.

An ad hoc network can be quickly deployed and provide much needed communications. Ad hoc networks will allow people to exchange data in the field or in a class room without using any network structure except the one they create by simply turning on their computers or PDAs, for example.

New applications for mobile ad hoc networks will continue to emerge and become an important part of the communication structure. Due to the lack of a fixed infrastructure, nodes must self-organize and reconfigure as they move, join or leave the network. All nodes could potentially be functionally identical and there may not be any natural hierarchy or central controller in the network. Many network-controlling functions are distributed among the nodes. Nodes are often powered by batteries and have limited communication and computation capabilities. The bandwidth of the system is usually limited. The distance between two nodes often exceeds the radio transmission range, and a transmission has to be relayed by other nodes before reaching its

destination. Consequently, a network has a multihop topology, and this topology changes as the nodes move around.

5 The Mobile Ad-Hoc Networks (MANET) working group of the Internet Engineering Task Force (IETF) has been actively evaluating and standardizing routing, including multicasting, protocols. Because the network topology changes arbitrarily as the nodes move, information is subject to becoming obsolete, and different nodes often have different views of the network, both in time (information may be outdated at some nodes but current at others) and in space (a node may only know the network topology in its neighborhood usually not far away from itself).

15 A routing protocol needs to adapt to frequent topology changes and with less accurate information. Because of these unique requirements, routing in these networks is very different from others. Gathering fresh information about the entire network is often costly and impractical. Many routing protocols are reactive (on-demand) protocols: they collect routing information only when necessary and to destinations they need routes to, and do not generally maintain unused routes after some period of time. This way the routing overhead is greatly reduced compared to pro-active protocols which maintain routes to all destinations at all times. It is important for a protocol to be adaptive. Ad Hoc on Demand Distance Vector (AODV), Dynamic Source Routing (DSR) and Temporally Ordered Routing Algorithm (TORA) are representative of on-demand routing protocols presented at the MANET working group.

30 Examples of other various routing protocols include Destination-Sequenced Distance Vector (DSDV) routing which is disclosed in U.S. Patent No. 5,412,654 to Perkins, and Zone Routing Protocol (ZRP) which is disclosed in U.S. Patent No. 6,304,556 to Haas. ZRP is a hybrid protocol using both

proactive and reactive approaches based upon distance from a source node.

These conventional routing protocols use a best effort approach in selecting a route from the source node to the destination node. Typically, the number of hops is the main criteria (metric) in such a best effort approach. In other words, the route with the least amount of hops is selected as the transmission route.

Quality-of-service (QoS) routing in mobile ad hoc networks is gaining interest. To provide quality-of-service, the protocol needs not only to find a route but also to secure the resources along the route. Because of the limited, shared bandwidth of the network, and lack of central controller which can account for and control these limited resources, nodes must negotiate with each other to manage the resources required for QoS routes. This is further complicated by frequent topology changes. Due to these constraints, QoS routing is more demanding than best-effort routing.

Some examples of QoS routing approaches are set forth by Chenxi Zhu in the publication entitled "Medium Access Control and Quality-of-Service Routing for Mobile Ad Hoc Networks," 2001, and by M. Mirhakkak et al. in the publication entitled "Dynamic Quality-of-Service for Mobile Ad Hoc Networks," MITRE Corp., 2000. Zhu discusses establishing bandwidth guaranteed QoS routes in small networks whose topologies change at a low to medium rate. Mirhakkak et al. are concerned with resource reservation requests which specify a range of QoS values while the network makes a commitment to provide service within this range.

The performance and reliability of reactive routing protocols such as DSR and AODV is a concern in the field of mobile ad-hoc networks. Conventional DSR finds multiple paths, but it only uses a single path for the transmission of message data/mission traffic. The remaining paths can be

backups, but are often stale when needed. Conventionally, AODV finds a single path. Failure of the primary path often results in significant delay because mission traffic must wait on a new route discovery process to complete.

5

Summary of the Invention

In view of the foregoing background, it is therefore an object of the present invention to provide more reliable and more timely routes with less traffic load concentration in a mobile ad hoc network.

10 This and other objects, features, and advantages in accordance with the present invention are provided by a method for routing message data from a source node to a destination node in a mobile ad hoc network. The mobile ad-hoc network including a plurality of intermediate mobile nodes between the source node and the destination node, and a plurality of wireless communication links connecting the nodes together. The method includes, at the source node, discovering routing to the destination node, ranking discovered routes according to at least one link metric, and distributing the message data 20 to the destination node along a plurality of the discovered routes based upon the ranking. The link metric may include a measurement of link delay, link capacity, link available capacity, and/or link reliability.

25 Distributing the message data may include distributing duplicate message data along the plurality of discovered routes, and/or distributing different portions of the message data along each of the plurality of discovered routes. Also, ranking the discovered routes preferably includes storing route entries in a route cache with each route entry corresponding to one of the discovered routes. Each route entry may include a metric value, a usage factor indicating how much of the message data should be distributed to the corresponding route, and a timer value for expiring the

discovered route. The route entries may be repeated in the route cache for different classes of message data, such as delay sensitive message data and large volume message data.

The method may also include, at the source node, intermediate nodes and/or the destination node, detecting whether the node can continue to support the route request and, if not, generating an error notification for a failed route. Upon receiving the error notification, the source node will purge the failed route from the discovered routes, and preferably distribute message data among remaining discovered routes. Additionally, the source node may transmit a second route request to discover new routing to the destination node when the number of discovered routes falls below a threshold.

Objects, features, and advantages in accordance with the present invention are also provided by a mobile ad hoc network including a plurality of mobile nodes, and a plurality of wireless communication links connecting the plurality of mobile nodes together. Each mobile node comprises a communications device to wirelessly communicate with other nodes of the plurality of nodes via the wireless communication links, and a controller to route communications via the communications device. The controller includes a route discovery unit to discover routing to a destination node, a route ranking unit to rank discovered routes according to at least one link metric, and a message data distribution unit to distribute the message data to the destination node along a plurality of the discovered routes based upon the rank. The link metric is preferably a measurement of link delay, link capacity, link available capacity, and/or link reliability.

The message data distribution unit may distribute duplicate message data along the plurality of discovered routes and/or different portions of the message data along each of the plurality of discovered routes. The route ranking unit comprises a route cache for storing route entries each

corresponding to one of the discovered routes. Again, each route entry may include a metric value, a usage factor indicating how much of the message data should be distributed to the corresponding route, and a timer value for expiring the discovered route. Also, the ranking unit may repeat route entries in the route cache for different classes of message data, such as delay sensitive message data and large volume message data.

Also, the controller may include a route request processing unit to determine whether the node can continue to support the route request and, if not, to generate an error notification for a failed route. The route ranking unit purges failed routes from the discovered routes, and the route discovery unit discovers new routing to the destination node when the number of discovered routes falls below a threshold.

Brief Description of the Drawings

FIGs. 1-4 are schematic diagrams of a mobile ad hoc network including in accordance with the present invention.

FIG. 5 is a flowchart illustrating the method steps for in a mobile ad hoc network in accordance with the present invention.

FIG. 6 is a schematic diagram illustrating a router of a node in accordance with the network of the present invention.

FIG. 7 is a schematic diagram illustrating the details of the controller of the router in FIG. 6.

Detailed Description of the Preferred Embodiments

The present invention will now be described more fully hereinafter with reference to the accompanying drawings, in which preferred embodiments of the invention are shown. This invention may, however, be embodied in many different forms and should not be construed as limited to the

embodiments set forth herein. Rather, these embodiments are provided so that this disclosure will be thorough and complete, and will fully convey the scope of the invention to those skilled in the art. Like numbers refer to like elements throughout, and prime notation is used to indicate similar elements in alternative embodiments.

As will be appreciated by those skilled in the art, portions of the present invention may be embodied as a method, data processing system, or computer program product.

Accordingly, these portions of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment, or an embodiment combining software and hardware aspects. Furthermore, portions of the present invention may be a computer program product on a computer-usable storage medium having computer readable program code on the medium. Any suitable computer readable medium may be utilized including, but not limited to, static and dynamic storage devices, hard disks, optical storage devices, and magnetic storage devices.

The present invention is described below with reference to flowchart illustrations of methods, systems, and computer program products according to an embodiment of the invention. It will be understood that blocks of the illustrations, and combinations of blocks in the illustrations, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, implement the functions specified in the block or blocks.

These computer program instructions may also be stored in a computer-readable memory that can direct a

computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory result in an article of manufacture including instructions which implement the
5 function specified in the flowchart block or blocks. The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer
10 implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions specified in the flowchart block or blocks.

Referring initially to FIGs. 1-5, a method for in a
15 mobile ad hoc network 20 will now be described. The network 20 includes a plurality of mobile nodes 30 including the source node 1 and the destination node 4 with intermediate nodes 2, 3 and 5 therebetween. The nodes 30, such as laptop computers, personal digital assistants (PDAs) or mobile
20 phones, are connected by wireless communication links 32 as would be appreciated by the skilled artisan. The method begins (block 100) and includes transmitting a route request RREQ from the source node 1 to discover routing to the destination node 4, as indicated at block 102 in FIG. 5. More
25 specifically, when a new route is needed to a given destination node 4, the source node 1 broadcasts the RREQ packet to the destination node. This RREQ packet is similar to the conventional RREQ packet used in a protocol such as DSR or AODV, but may include a link/route metric, as will be
30 discussed in more detail below. The conventional RREQ broadcast is used for "best effort" service. The method of the present invention may still follow the conventional procedures established by the protocol for best effort service.

Furthermore, at block 104, the method includes each intermediate node 2, 3 and 5 determining whether the node can support the route request RREQ. If the node cannot support the particular request RREQ, then the request is denied or simply not forwarded by the node (block 106). If the node, for example node 3, can support the particular request RREQ, then the node forwards the route request to other intermediate nodes 2 and 5 (block 108), and may temporarily reserve node resources for that route request. Intermediate nodes 2 and 5 also must determine whether they can support the route request RREQ forwarded from node 3. If so, the route request RREQ is then forwarded to the destination node 4.

The destination node 4, upon receiving the route request RREQ, generates a reply RREP to the source node 1 for each discovered route (block 110). In other words, the destination node 4 may have received the forwarded route request RREQ from any of various possible routes including, for example, 1-2-4 or 1-3-5-4. A reply RREP is generated in each case. At block 112, the source node 1 ranks the discovered routes according to one or more link metrics. The link metric is preferably a measurement of link delay, link capacity, link available capacity, and/or link reliability as will be discussed below.

Ranking the discovered routes preferably includes storing route entries in a route cache or table. Each of the route entries corresponds to one of the discovered routes. Also, each route entry may include a metric value, a usage factor indicating how much of the message data should be distributed to the corresponding route, and a timer for expiring the discovered route. Some or all of the route entries may be repeated for different classes of message data/traffic, such as, delay sensitive traffic and/or large volume traffic, for example.

At block 114, the source node 1 then selects a plurality of routes to the destination node 4, and, at block 116, the source node may transmit route confirmations CONF to intermediate nodes on the selected route. This is to confirm the use of the resources on the selected route in a system using a routing approach other than a best effort approach, such as quality of service routing.

Also, at block 118, the source node 1 distributes the message data to the destination node 4 along the plurality of discovered routes, such as routes 1-2-4 and 1-3-5-4, for example. Distributing the message data may include distributing duplicate message data along the plurality of discovered routes for improved reliability, and/or distributing different portions of the message data along each of the plurality of discovered routes for improved timeliness. As used herein, "message data" is intended to include any data that may be sent between nodes in the mobile ad-hoc network, including (but not limited to) video data, audio data, alphanumeric data, etc., as would be appreciated by the skilled artisan.

At block 120, the intermediate nodes 2, 3 and 5, and/or the destination node 4, may detect at any time whether the node can continue to support the route request RREQ. If the node cannot continue to support the request RREQ, then the node generates an error notification RERR to the source node 1. Here, (block 122) the source node 1 may purge the failed route, upon receiving the error notification RERR, while maintaining the use of the other discovered routes. However, if the number of discovered routes falls below a certain threshold, then the source node 1 may again transmit a route request RREQ to discover new routing to the destination node 4 (block 124).

The described method can be applied to any type of On-Demand or Reactive routing protocol, such as Dynamic Source

Routing (DSR) or Ad-Hoc On-Demand Distance Vector (AODV) routing, or to any hybrid proactive/reactive protocol, such as Zone Routing Protocol (ZRP), as would be appreciated by the skilled artisan. The described procedures are easily applied to the DSR protocol. The conventional DSR message types RREQ, RREP, RREP, RREP are defined as optional packet types, and can be used as defined for the conventional operation of the protocol to support "best effort" traffic in a backwards compatibility mode. New optional packet types may be defined to support the measurement of certain link metrics. Definition of the required header fields for these types is straightforward based on the functions defined above.

As described, the performance of reactive ad hoc routing protocols such as DSR and AODV are improved to provide more reliable and more timely routes with less traffic load concentration. Routes are discovered via route-request and route-reply scheme. It should be noted that conventional DSR finds multiple paths, but it only uses a single path for message data transmission while remaining paths can be backups, but are often stale when needed. Conventionally, AODV finds a single path. Failure of primary path often results in significant delay because mission traffic must wait on a new route discovery process to complete.

As discussed, the present invention provides the discovery and use of multiple routes with appropriate metrics for each to the destination. Discovered paths are ranked according to metric and/or traffic class. Routing cache/table is tailored to utilize up to n routes simultaneously per destination (may be used either to distribute the load or for reliability). Traffic is distributed across multiple routes according to rank and value indicated by the path metrics. Usage of n paths allows simultaneous maintenance of n paths while failure of a single path can leave n-1 paths still operational.

The route cache/table includes destination indexes into a group of routes (either source routes for DSR or next hop routes for AODV). Each route entry has a metric value, a usage factor (e.g. between 0 and 1) that indicates which
5 fraction of the traffic load should use that route, and a timer for expiring the route. In use, one of up to n routes is selected according to any approach that allows the load to be shared (such as random selection or a round robin servicing according to usage factor). Route entries may also be
10 repeated for each class of traffic (e.g. using different metrics and ranking for each class).

Regarding route failure: Nodes discovering route failure generate route error packets (RERR) and send back to the traffic source. Traffic source purges failed routes from
15 cache/table. Traffic source continues to use the n-1 "good" routes. If other routes are available (with satisfactory metric) then they are added to the table to obtain n routes. Usage factors are recalculated. If a failure causes the number of routes in use to fall below a predefined threshold,
20 a new route discovery is initiated to find additional routes.

Regarding the use of metrics: route requests and route replies are modified to accommodate improved metrics. DSR and AODV currently use only hop count which can still be used if desired. Improved metrics can include a variety of
25 measures such as link delay, link capacity, link available capacity, link reliability. More important metrics may depend on the traffic class such as: Delay sensitive traffic may require routes to be rank-ordered based on the delay metric; Large volume traffic may require routes to be rank-ordered
30 based on the capacity metric.

The usage factor for each route as may be calculated as: $\text{UsageFactor}(\text{Rtek}) = F(\text{metric}_k, f(\text{metrics for all paths used}))$. For large volume traffic, bottleneck capacity may be used as a metric in evaluating routes - i.e.,

UsageFactor(Rtex)=BottleneckCapacity_k/(sum over all routes
(BottleneckCapacity for each route)).

The invention as specifically applied to DSR will be described. DSR naturally collects multiple routes but only
5 one is used. Loop freedom is maintained with multiple routes because the source route indicates the full path (loops are easily detected). DSR is modified to collect appropriate metrics for each link on the source route. The DSR route cache is modified to reflect the use of multiple routes
10 simultaneously for multiple classes of traffic using the route usage factors. Ranking of routes for each traffic class is implemented in the route cache and the top n routes for each traffic class are selected. Because source routes have been collected, the ranking process may include a factor such as
15 disjointness of paths when determining which n paths are the most desirable (disjoint paths would provide a set of paths that are less likely to be simultaneously disrupted due to link/node failures and mobility). Usage factors for top n routes for each traffic class would be calculated based upon
20 the appropriate path metrics as indicated above. Traffic distribution among several paths is done only at the source node. A timer to expire stale routing cache entries may be used (i.e. optional). A higher reliability traffic class which sends duplicate traffic over multiple routes for higher
25 reliability may be implemented. Route failure recovery is implemented as discussed previously. This will naturally discover any stale routes in the top n routes that are being actively used.

The invention as specifically applied to AODV will
30 be described. AODV collects only a single route as currently defined with path hop count as the metric. Thus, AODV should be modified to collect multiple routes with the appropriate metrics. Care should be taken with AODV to insure loop-free routes since it builds next hop routing tables. For example,

the following conservative approach will work: Modify AODV route table to reflect the use of multiple routes, the ranking of the top n routes, and the usage factors for top n routes based upon the appropriate path metrics; select one of the top n routes the "best route" - to maintain loop-freedom it will be selected in AODV as the freshest minimum hop count route (the lesser routes included in the top n can have up to 1 greater hop count); Intermediate nodes return RREPs designating only their choice for "best route" while the destination can return multiple RREPs for paths through different neighbors; Selection of one of n possible paths occurs only at the source node - each intermediate forwarding node always forwards a packet along its "best route". When a node finds a route that is fresher than another route already in use it does not have to discard the other route unless it fails or hop count is too high. Higher reliability traffic class which sends duplicate traffic over multiple routes for higher reliability may be implemented. Route failure results in generation of RERRs and triggering the failure recovery process as indicted above. Thus, as described, the present invention supports conventional ad hoc routing protocols such as DSR and AODV but significantly improves performance. Multiple routes can be used either for load balancing or increased probability of data traffic delivery. The use of multiple routes can provide more timely backup routes when path failures occur. In most cases failures will not disrupt data traffic other than the traffic temporarily lost on the failed link. Route discovery can be initiated before all routes have failed reducing the likelihood of having no routes, i.e. route discovery that is initiated after all routes are lost will be much rarer.

Referring now additionally to Figs. 6 and 7, a system aspect of the invention will be described. A mobile ad hoc network 20 includes a plurality of mobile nodes 30, and a

plurality of wireless communication links 32 connecting the plurality of mobile nodes together. Each mobile node includes a router 40 (FIG. 6) that has a communications device 42 to wirelessly and uni-directionally or bi-directionally
5 communicate with other nodes via the wireless communication links 32, and a controller 44 to route communications via the communications device 42. Also, a memory 46 may be included as part of the controller 44 or in connection with the controller.

10 As shown in FIG. 7, the controller 44 includes a route discovery unit 50 to transmit the route requests to other nodes to discover routing to a destination node. Also, a route request processing unit 52 determines whether the node can support a route request. Furthermore, the route request
15 processing unit 52 detects whether the node can continue to support the route request and, if not, generates the error notification RERR. A route ranking unit 54 ranks discovered routes according to one or more link metrics, and a message
20 data distribution unit 58 distributes the message data to the destination node along a plurality of the discovered routes based upon the rank. Again, the link metric is preferably a measurement of link delay, link capacity, link available capacity, and/or link reliability.

The message data distribution unit 58 may distribute
25 duplicate message data along the plurality of discovered routes and/or different portions of the message data along each of the plurality of discovered routes. The route ranking unit 54 comprises a route cache 56 for storing route entries each corresponding to one of the discovered routes. Again,
30 each route entry may include a metric value, a usage factor indicating how much of the message data should be distributed to the corresponding route, and a timer for expiring the discovered route. Also, the ranking unit 54 may repeat route entries in the route cache for different classes of message

data, such as delay sensitive message data and large volume message data. The route ranking unit 54 purges failed routes from the discovered routes, and the route discovery unit 50 discovers new routing to the destination node when the number
5 of discovered routes falls below a threshold.

CLAIMS

1. A method for routing message data from a source node to a destination node in a mobile ad hoc network comprising a plurality of intermediate mobile nodes between the source node and the destination node, and a plurality of wireless communication links connecting the nodes together, the method comprising:
- at the source node, discovering routing to the destination node;
 - at the source node, ranking discovered routes according to at least one link metric; and
 - at the source node, distributing the message data to the destination node along a plurality of the discovered routes based upon the ranking.
2. A method according to Claim 1 wherein ranking the discovered routes comprises storing route entries in a route cache, each route entry corresponding to one of the discovered routes.
3. A method according to Claim 2 wherein each route entry includes:
- a metric value for the link metric; and
 - a usage factor indicating how much of the message data should be distributed to the corresponding route.
4. A method according to Claim 3 wherein each route entry further includes a timer value for expiring the discovered route.
5. A method according to Claim 1 further comprising, at the source node, detecting failed routes and purging failed routes from the discovered routes.

6. A mobile ad hoc network comprising:
a plurality of mobile nodes; and
a plurality of wireless communication links
connecting the plurality of mobile nodes together;
5 each mobile node comprising
a communications device to wirelessly
communicate with other nodes of the plurality of
nodes via the wireless communication links, and
a controller to route communications via the
10 communications device, the controller comprising
a route discovery unit to discover
routing to a destination node,
a route ranking unit to rank discovered
routes according to at least one link metric,
15 and
a message data distribution unit to
distribute the message data to the destination
node along a plurality of the discovered
routes based upon the rank.

20 7. A network according to Claim 6 wherein the
route ranking unit comprises a route cache for storing route
entries each corresponding to one of the discovered routes.

25 8. A network according to Claim 7 wherein each
route entry includes:
a metric value for the link metric; and
a usage factor indicating how much of the message
data should be distributed to the corresponding route.

30 9. A network according to Claim 8 wherein each
route entry further includes a timer value for expiring the
discovered route.

10. A network according to Claim 6 wherein the link metric comprises a measurement of at least one of link delay, link capacity, link available capacity, and link reliability.

5

11. A network according to Claim 6 wherein the controller further comprises a route request processing unit to determine whether the node can continue to support the route request and, if not, to generate an error notification

10

for a failed route.

1/4




FIG. 1.




FIG. 2.

2/4




FIG. 3.




FIG. 4.

3/4




FIG. 5.

4/4




FIG. 6.




FIG. 7.